

# Conflict diagnostics for evidence synthesis in a multiple testing framework

Anne M. Presanis, David Ohlssen, Kai Cui,  
Magdalena Rosinska, Daniela De Angelis

February 24, 2017

*Medical Research Council Biostatistics Unit, University of Cambridge, U.K.*  
*Novartis Pharmaceuticals Corporation, East Hanover, NJ, U.S.A.*  
*Department of Epidemiology, National Institute of Public Health,*  
*National Institute of Hygiene, Warsaw, Poland*

e-mail: anne.presanis@mrc-bsu.cam.ac.uk

## Abstract

Evidence synthesis models that combine multiple datasets of varying design, to estimate quantities that cannot be directly observed, require the formulation of complex probabilistic models that can be expressed as graphical models. An assessment of whether the different datasets synthesised contribute information that is consistent with each other (and in a Bayesian context, with the prior distribution) is a crucial component of the model criticism process. However, a systematic assessment of conflict in evidence syntheses suffers from the multiple testing problem, through testing for conflict at multiple locations in a model. We demonstrate how conflict diagnostics can be employed throughout a graphical model while accounting for the multiple hypothesis tests of no conflict at each location in the graph. The method is illustrated by a network meta-analysis to estimate treatment effects in smoking cessation programs and an evidence synthesis to estimate HIV prevalence in Poland.

**KEYWORDS:** *Conflict; evidence synthesis; graphical models; model criticism; multiple testing; network meta-analysis.*

## 1 Introduction

Evidence synthesis refers to the use of complex statistical models that combine multiple, disparate and imperfect sources of evidence to estimate quantities on which direct information is unavailable or inadequate (e.g. Ades and Sutton, 2006; Welton et al., 2012; De Angelis et al.,

2014). More formally, our goal is to estimate  $K$  *basic* parameters  $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)$  given a collection of  $N$  independent data sources  $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)$ , where each  $\mathbf{y}_i, i \in 1, \dots, N$  may be a vector or array of data points. Each  $\mathbf{y}_i$  provides information on a *functional* parameter  $\psi_i$  (or potentially a vector of functions  $\boldsymbol{\psi}_i$ ). When  $\psi_i = \theta_k$  is the identity function, the data  $\mathbf{y}_i$  are said to *directly* inform  $\theta_k$ . Otherwise,  $\psi_i = \psi_i(\boldsymbol{\theta})$  is a function of multiple parameters in  $\boldsymbol{\theta}$ : the  $\mathbf{y}_i$  therefore only provide *indirect* information on these parameters. Given the conditional independence of the datasets  $\mathbf{y}_i$ , the likelihood is  $L(\boldsymbol{\theta}; \mathbf{y}) = \prod_{i=1}^N L_i(\psi_i(\boldsymbol{\theta}); \mathbf{y}_i)$ , where  $L_i(\psi_i(\boldsymbol{\theta}); \mathbf{y}_i)$  is the likelihood contribution of  $\mathbf{y}_i$  to the basic parameters  $\boldsymbol{\theta}$ . In a Bayesian context, for a prior distribution  $p(\boldsymbol{\theta})$ , the posterior distribution  $p(\boldsymbol{\theta} \mid \mathbf{y}) \propto p(\boldsymbol{\theta})L(\boldsymbol{\theta}; \mathbf{y})$  summarises all information, direct and indirect, on  $\boldsymbol{\theta}$ . Let  $\boldsymbol{\psi} = (\psi_1, \dots, \psi_N)$  be the set of functional parameters informed by data and  $\boldsymbol{\phi} = \{\boldsymbol{\theta}, \boldsymbol{\psi}\}$  be the set of all unknown quantities, whether basic or functional. Such evidence synthesis models may be represented by directed acyclic graphs (DAGs), encoding the conditional independence assumptions made (Lauritzen, 1996). Denote a DAG by  $\mathcal{G}(\mathbf{V}, \mathbf{E})$ , where  $\mathbf{V} = \{\boldsymbol{\phi}, \mathbf{y}\}$  is the set of nodes in the graph representing either known or unknown quantities; and  $\mathbf{E}$  is the set of directed edges representing dependencies between nodes. Each ‘child’ node is independent of its ‘siblings’ conditional on their direct ‘parents’. The joint distribution of all nodes  $\mathbf{V}$  is the product of the conditional distributions of each node given its direct parents. An example DAG of an evidence synthesis model is given in Figure 1(i). Circles denote unknown quantities: either basic parameters  $\boldsymbol{\theta}$  that are ‘founder’ nodes at the top of a DAG having a prior distribution (double circles); or functional parameters  $\boldsymbol{\psi}$ . Squares denote observed quantities, solid arrows represent stochastic distributional relationships, and dashed arrows represent deterministic functional relationships. This DAG could be extended to more complex hierarchical priors and models, where repetition over variables may be represented by ‘plates’, rounded rectangles around the repeated nodes, with the label giving the range of repetition. In general, the set  $\mathbf{V}$  may be larger than just the set of basic and functional parameters, including also other intermediate nodes in the DAG, for example unit-level parameters in a hierarchical model. For brevity, from here on we will abbreviate any DAG to the notation  $\mathcal{G}(\boldsymbol{\phi}, \mathbf{y})$ .

With increased computational power, models of the form of  $\mathcal{G}(\boldsymbol{\phi}, \mathbf{y})$  have proliferated, requiring also the development of model criticism tools adapted to the challenges of evidence synthesis. In a Bayesian framework, any of the prior distribution, the assumed form of the likelihood and structural and functional assumptions may conflict with the observed data or with each other. To assess the consistency of each of these components, various mixed- or posterior-predictive checks have been proposed. In particular, the “conflict p-value” (Marshall and Spiegelhalter, 2007; Gåsemyr and Natvig, 2009; Presanis et al., 2013; Gåsemyr, 2016) is a diagnostic calculated by splitting  $\mathcal{G}(\boldsymbol{\phi}, \mathbf{y})$  into two independent sub-graphs (“partitions”) at a particular “separator” node  $\phi$ , to measure the consistency of the information provided by each partition about the node. The separator  $\phi$  could either be a founder node, i.e. a basic parameter, or a node internal to the DAG, and is split into two copies  $\phi_a$  and  $\phi_b$ , one in each partition (a “node-split”, see Figure 1(ii,iii)). Suppose that partition  $a$  contains the data vector  $\mathbf{y}_a$  and provides inference resulting in a posterior distribution  $p(\phi_a \mid \mathbf{y}_a)$ , and that similarly partition  $b$  results in  $p(\phi_b \mid \mathbf{y}_b)$ . The aim is to assess the null hypothesis that  $\phi_a = \phi_b$ . For  $\phi$  taking discrete values, we can directly evaluate  $p(\phi_a = \phi_b \mid \mathbf{y}_a, \mathbf{y}_b)$ . If the support of  $\phi$  is continuous, we consider the posterior probability of  $\delta = h(\phi_a) - h(\phi_b)$ , where  $h(\cdot)$  is a function that transforms  $\phi$  to a scale for which a uniform prior is appropriate. The two-sided “conflict p-value” is defined as  $c = 2 \times \min \{\Pr\{p_\delta(\delta \mid \mathbf{y}_a, \mathbf{y}_b) < p_\delta(0 \mid \mathbf{y}_a, \mathbf{y}_b)\}, 1 - \Pr\{p_\delta(\delta \mid \mathbf{y}_a, \mathbf{y}_b) < p_\delta(0 \mid \mathbf{y}_a, \mathbf{y}_b)\}\}$ ,




Figure 1: (i) Example DAG  $\mathcal{G}(\mathbf{V}, \mathbf{E})$  showing a generic evidence synthesis. (ii) & (iii) Example node-split at separator node  $\phi$ : (ii) original model  $\mathcal{G}(\phi, \mathbf{y})$ ; (iii) node-split model. In (ii): the data  $\mathbf{y} = \{\mathbf{y}_\phi, \mathbf{y}_{\bar{\phi}}\}$  comprise data  $\mathbf{y}_\phi$  that are direct descendents of  $\phi$ ; and the remaining data  $\mathbf{y}_{\bar{\phi}}$ . In (iii): when splitting  $\mathcal{G}(\phi, \mathbf{y})$  into partitions  $a$  and  $b$ , the data vector  $\mathbf{y}_\phi$  is split into  $\mathbf{y}_{a,\phi}$  and  $\mathbf{y}_{b,\phi}$ , whereas  $\mathbf{y}_{\bar{\phi}}$  remains only in partition  $a$ . The partition  $a$  data are therefore  $\mathbf{y}_a = \{\mathbf{y}_{a,\phi}, \mathbf{y}_{\bar{\phi}}\}$  and the partition  $b$  data are  $\mathbf{y}_b = \mathbf{y}_{b,\phi}$ .

where  $p_\delta$  is the posterior density of the difference  $\delta$ , so that the smaller  $c$  is, the greater the conflict between the two partitions. G  sem  r and Natvig (2009) and Presanis et al. (2013) demonstrate how  $c$  may be evaluated in different contexts, including both one- and two-sided hypothesis tests, and G  sem  r (2016) demonstrates the uniformity of the conflict p-value in a wide range of models.

The conflict p-value may be used in a targeted manner, searching for conflict at particular nodes in a DAG. However, in complex evidence syntheses, often the location of potential conflict may be unclear. A systematic assessment of conflict throughout a DAG is then required to locate problem areas (e.g. Krahn et al., 2013). Such systematic assessment, however, suffers from the multiple testing problem, either through testing for conflict at each node in  $\mathcal{G}(\phi, \mathbf{y})$  or through the separation of  $\mathcal{G}(\phi, \mathbf{y})$  into more than two partitions to simultaneously test for conflict between each pair-wise partition. We account for these multiple tests by adopting the general hypothesis testing framework of Hothorn et al. (2008); Bretz et al. (2011), allowing for simultaneous multiple hypotheses in a parametric setting. They propose different possible tests to account for multiplicity: we concentrate here on maximum-T type tests.

In section 2, we introduce the particular evidence synthesis models that motivate our work on systematic conflict assessment, a network meta-analysis and a model for estimating HIV prevalence. Section 3 describes the methods we use to test for conflict and account for the multiple tests we perform. We apply these methods to our examples in Section 4 and end with a discussion in Section 5.

## 2 Motivating examples

### 2.1 Network meta-analysis

Network meta-analysis (NMA) is a specific type of evidence synthesis (Salanti, 2012), that generalises meta-analysis from the synthesis of studies measuring a treatment effect (e.g. of treatment B versus treatment A in a randomised clinical trial), to the synthesis of data on more than two treatment arms. The studies included in the NMA may not all measure the same treatment effects, but each study provides data on at least two of the treatments. For example, considering a set of treatments  $\{A, B, C, D\}$ , the network of trials may consist of studies of different ‘‘designs’’, i.e. with different subsets of the treatments included in each trial (Jackson et al., 2014), such as  $\{ABC, ABD, BD, CD\}$ . As with meta-analysis, NMA models can be implemented in either a two-stage or single-stage approach, as described more comprehensively elsewhere (Salanti, 2012; Jackson et al., 2014).

Here we concentrate on a single-stage approach, where the original data  $Y_{di}^J$  for each treatment  $J$  of each study  $i$  of design  $d$  are available. A full likelihood model specifies

$$Y_{di}^J \sim f(p_{di}^J \mid w_{di}^J)$$

for some distribution  $f(\cdot)$  and treatment outcome  $p_{di}^J$  with associated information  $w_{di}^J$ . For example, if the data are the numbers of events out of total numbers at risk of the event, then  $w_{di}^J$  might be the denominator for treatment  $J$ . Then we might assume the data

are realisations of a Binomial random variable,  $Y_{di}^J \sim \text{Bin}(w_{di}^J, p_{di}^J)$ , where the proportion  $p_{di}^J$  is a function of a study-specific baseline  $\alpha_{di}$  representing a design/study-specific baseline treatment  $B_d$  and a study-specific treatment contrast (log odds ratio)  $\mu_{di}^{B_d J}$ , through a logistic model,  $\text{logit}(p_{di}^J) = \alpha_{di} + \mu_{di}^{B_d J}$ . Note that the intercept  $\alpha_{di} = \text{logit}(p_{di}^{B_d})$ .

To complete the model specification requires parameterisation of the treatment effects  $\mu_{di}^{AJ}$ . A common effect model, for a network-wide reference treatment  $A$ , is given by

$$\mu_{di}^{AJ} = \eta^{AJ} \quad (1)$$

for each  $J \neq A$ , i.e. assumes that all studies of all designs measure the same treatment effects. The  $\eta^{AJ}$  are basic parameters, of which there are the number of treatments in the network minus 1, representing the relative effectiveness of treatment  $J$  compared to the network baseline treatment  $A$ . All other contrasts  $\eta^{JK}$ ,  $J, K \neq A$  are functional parameters, defined by assuming a set of *consistency* equations  $\eta^{JK} = \eta^{AK} - \eta^{AJ}$  for each  $J, K \neq A$ . These equations define a transitivity property of the treatment effects.

The extension to a random-effects model, still under the consistency assumption, implies

$$\mu_{di}^{AJ} = \eta^{AJ} + \beta_{di}^{AJ} \quad (2)$$

where usually the random effects  $\beta_{di}^{AJ}$ , reflecting between-study heterogeneity, are assumed normally distributed around 0, with a covariance structure defined as a square matrix  $\Sigma_\beta$  such that all entries on the leading diagonal are  $\sigma_\beta^2$  and all remaining entries are  $\sigma_\beta^2/2$  (Salanti, 2012; Jackson et al., 2015). Figure .1 of the Supplementary Material shows the DAG structure of both the common and the random effects models for a full likelihood setting where the outcome is binomial. The set of *basic* parameters is denoted  $\boldsymbol{\eta}_b = (\eta^{AJ})_{J \neq A}$  and the corresponding set of *functional* parameters is denoted  $\boldsymbol{\eta}_f = (\eta^{JK} = \eta^{AK} - \eta^{AJ})_{J, K \neq A}$ . Note that the common-effect model is a special case of the random-effects model.

In the Bayesian paradigm, we specify prior distributions for the basic parameters  $\boldsymbol{\eta}_b$ , the (nuisance) study-specific baselines  $\alpha_{di}$ , and in the case of the random treatment effects model, the common standard deviation parameter  $\sigma_\beta$  in terms of which the variance-covariance matrix  $\Sigma_\beta$  is defined. Note that any change in parameterisation of the model, for example changing treatment labels, will affect the joint prior distribution, making invariance challenging or even impossible in a Bayesian setting.

**A smoking cessation example** Dias et al. (2010), amongst many others (Lu and Ades, 2006; Higgins et al., 2012; Jackson et al., 2015), considered an NMA of studies of smoking cessation. The network consists of 24 studies of 8 different designs, including 2 three-arm trials. Four smoking cessation counselling programs are compared (Figure 2): A = no intervention; B = self-help; C = individual counselling; D = group counselling. The data (Table .1 of the Supplementary Material) are the number of individuals out of those participating who have successfully ceased to smoke at 6-12 months after enrollment. Here we fit the common- and random-effect models under a consistency assumption and diffuse priors: Normal centred on 0, with standard deviation of 10 on the log-odds scale for  $\boldsymbol{\eta}_b$  and  $\alpha_{di}$ ; and Uniform(0, 5) for  $\sigma_\beta$ . We find (Table .2 of the Supplementary Material) that the deviance information criterion (*DIC*, Spiegelhalter et al. (2002)) prefers the random-effect model, suggesting the random effects are necessary to explain the heterogeneity in the network. Note that the estimates of the treatment effects from the random-effect model are both somewhat different and more

uncertain than those from the common-effect model, agreeing with estimates found by other authors, including Dias et al. (2010). Note also that the posterior expected deviance for the random-effect model,  $\mathbb{E}_{\theta|y}(D) = 54$ , is slightly larger than the number of data points (50), suggesting still some lack of fit to the data (Spiegelhalter et al., 2002).

**A single node-split model** This residual lack of fit and the potential in NMA in general for variability between groups of direct and indirect information from multiple studies that is excess to between-study heterogeneity (“inconsistency”, Lu and Ades (2006)) has motivated various approaches to the detection and resolution of inconsistency (Lumley, 2002; Lu and Ades, 2006; Dias et al., 2010; Higgins et al., 2012; White et al., 2012; Jackson et al., 2014). Dias et al. (2010) apply the idea of node-splitting, based on Marshall and Spiegelhalter (2007), to the NMA context, splitting a single mean treatment effect  $\eta^{JK}$  in the random effects (consistency) model (2). The authors partition a DAG of the model into *direct* evidence from studies directly comparing  $J$  and  $K$  versus *indirect* evidence from all remaining studies. Specifically, for any study  $i$  of design  $d$  that directly compares  $J$  and  $K$ , the study-specific treatment effect is expressed in terms of the direct treatment effect:  $\mu_{di}^{JK} = \eta_{Dir}^{JK} + \beta_{di}^{JK}$ ; whereas the indirect version of the treatment effect is estimated from the remaining studies via the consistency equation:  $\eta_{Ind}^{JK} = \eta^{AK} - \eta^{AJ}$ . The posterior distribution of the contrast (or the inconsistency parameter)  $\delta^{JK} = \eta_{Dir}^{JK} - \eta_{Ind}^{JK}$  is then examined to check posterior support for the null hypothesis  $\delta^{JK} = 0$ .

**Multiple node-splits** Although the single node-split approach in Dias et al. (2010) has been extended to automate the generation of different single node-splitting models for conflict assessment (van Valkenhoef et al., 2016), the simultaneous splitting of multiple nodes (equivalently multiple treatment comparisons) in a NMA has not yet been considered. In section 4.1.1, we investigate whether there is conflict in the smoking cessation network beyond heterogeneity, while accounting for the multiple tests for inconsistency we perform.

## 2.2 Generalised evidence synthesis

As further illustration of systematic conflict detection, we consider an evidence synthesis approach to estimating HIV prevalence in Poland, among the exposure group of men who have sex with men (MSM) (Rosinska et al., 2016). The data aggregated to the national level are given in Table .3 of the Supplementary Material. There are three basic parameters to be estimated: the proportion of the male population who are MSM,  $\rho$ ; the prevalence of HIV infection in the MSM group,  $\pi$ ; and the proportion of those infected who are diagnosed,  $\kappa$  (Figure 3(a)).

**Likelihood** The total population of Poland,  $N = 15,749,944$ , is considered fixed. The remaining 5 data points  $y_1, \dots, y_5$  directly inform, respectively:  $\rho$ ; prevalence of diagnosed infection  $\pi\kappa$ ; prevalence of undiagnosed infection  $\pi(1 - \kappa)$ ; and lower ( $D_L$ ) and upper ( $D_U$ ) bounds for the number of diagnosed infections  $D = N\rho\pi\kappa$  (Figure 3(a), Table .3 of the Supplementary Material). These data are modelled independently as either Binomial samples ( $y_1, y_2, y_3$ ) or Poisson counts ( $y_4, y_5$ ).

**Priors** The number diagnosed  $D$  is constrained *a priori* to lie between the stochastic bounds  $D_L$  and  $D_U$ , which in turn are given vague log-normal priors. Since  $D$  is already defined as a function of the basic parameters, the constraint is implemented via introduction of an auxiliary Bernoulli datum of observed value 1, with probability parameter given by a functional parameter  $c = Pr(D_L \leq D \leq D_U)$  (Figure 3(a)). The basic parameters  $\rho, \pi$  and  $\kappa$  are given independent uniform prior distributions on  $[0, 1]$ .

**Exploratory model criticism** This initial analysis reveals a lack of fit to some of the data (Table .3 of the Supplementary Material), with particularly high posterior mean deviances for the data informing  $\rho$  and  $\pi\kappa$ . This lack of fit in turn may suggest the existence of conflict in the DAG (Spiegelhalter et al., 2002). In Rosinska et al. (2016), conflict between evidence sources was not directly considered or formally measured, instead resolving the lack of fit by modelling potential biases in the data in a series of sensitivity analyses. By contrast, in Section 4.2 we systematically assess the consistency of evidence coming from the prior model and from each likelihood contribution, by splitting the DAG at each functional parameter (Figure 3(b)).

## 3 Methods

### 3.1 Defining the hypothesis tests of conflict

In Presanis et al. (2013) and Section 1, we considered splitting a DAG  $\mathcal{G}(\phi, \mathbf{y})$  into two independent partitions  $\mathcal{G}(\phi_a, \mathbf{y}_a)$  and  $\mathcal{G}(\phi_b, \mathbf{y}_b)$ . Generalising this set-up, suppose that  $\mathcal{G}(\phi, \mathbf{y})$  is now partitioned into  $Q$  independent sub-graphs,  $\mathcal{G}_1(\phi_1, \mathbf{y}_1), \dots, \mathcal{G}_Q(\phi_Q, \mathbf{y}_Q)$ , where each disjoint subset of the data  $\mathbf{y}_q, q \in 1, \dots, Q$  is chosen to identify part of the basic parameter space  $\theta_q = (\theta_{q1}, \dots, \theta_{qb_q})$ , where  $b_q$  is the number of basic parameters in partition  $q$ . Note that  $\theta_q \subset \phi_q$  for each  $q \in 1, \dots, Q$ , whereas the complementary subset  $\phi_q \setminus \theta_q$  consists of functional and other non-basic parameters. To test the consistency of the information provided by each partition about a set of  $J$  separator nodes  $(\phi_1^{(s)}, \dots, \phi_J^{(s)}) \subseteq \phi$  from the original model, a set of contrasts  $\delta_j = (\delta_{j1}, \dots, \delta_{jC_j})$  is formed for each  $j \in 1, \dots, J$ , one contrast for each pair of partitions in which  $\phi_j$  appears. There is therefore a maximum of  $\binom{Q}{2}$  possible contrasts for each separator node, i.e.  $C_j \leq \binom{Q}{2}$ . Each contrast  $\delta_{jc}$  is defined as

$$\delta_{jc} = h_j(\phi_{jq_A} \mid \mathbf{y}_A) - h_j(\phi_{jq_B} \mid \mathbf{y}_B)$$

for the pair of partitions  $c = \{q_A, q_B\}$  and node-split copies  $\{\phi_{jq_A}, \phi_{jq_B}\}$ . The functions  $h_j(\cdot)$  are functions that transform the separator nodes  $\{\phi_{jq_A}, \phi_{jq_B}\}$  to an appropriate scale for a uniform (Jeffreys) prior to be applicable, if either is a founder node in either partition.

Denote the separator nodes in each partition by  $\phi_q^{(s)} = \{\phi_{jq}, j \in 1, \dots, m_q, q \in 1, \dots, Q\}$ , where  $m_q \leq J$  is the number of separator nodes in partition  $q$ . Writing these nodes as a stacked vector  $\phi_S = (\phi_1^{(s)}, \dots, \phi_Q^{(s)}) = (\phi_{11}, \dots, \phi_{m_1 1}, \phi_{12}, \dots, \phi_{m_2 2}, \dots, \phi_{1Q}, \dots, \phi_{m_Q Q})^T$ , and the transformed version as  $\phi_H = \mathbf{h}(\phi_S)$ , the total set of contrasts is

$$\Delta = (\delta_1, \dots, \delta_J)^T = \mathbf{C}_\Delta^T \phi_H$$

for an appropriate contrast matrix of 1s and 0s,  $\mathbf{C}_\Delta^T$ . Note that not every separator node necessarily appears in every partition, so although  $\phi_H$  has maximum length  $J \times Q$ , in practice, its length  $m = \sum_{q=1}^Q m_q \leq J \times Q$ . The contrast matrix  $\mathbf{C}_\Delta^T$  therefore has dimension  $p \times m$ , so that it maps from the space of the  $m$  separator nodes (including node-split copies) to that of the  $p = \sum_{j=1}^J C_j$  contrasts. A test for consistency of the information in each partition may be expressed as a test of the null hypothesis that

$$H_0 : \Delta = \mathbf{C}_\Delta^T \phi_H = \mathbf{0} \quad (3)$$

### 3.2 Asymptotic theory

Using standard asymptotic theory (Bernardo and Smith, 1994, see also derivation in Appendix B of the Supplementary Material), it can be shown that if the joint posterior distribution of all parameters  $\phi$  in all partitions is asymptotically multivariate normal (i.e. if the prior is flat enough relative to the likelihood), and if  $\frac{\partial \Delta(\phi)}{\partial \phi} = \mathbf{C}_\Delta^T$  is non-singular with continuous entries, then the posterior mean of  $\Delta$  is  $\bar{\Delta} = \mathbf{C}_\Delta^T \bar{\phi}_H \stackrel{a}{\approx} \mathbf{C}_\Delta^T \hat{\phi}_H$  and the posterior variance-covariance matrix of  $\Delta$  is  $\mathbf{S}_\Delta \stackrel{a}{\approx} \mathbf{C}_\Delta^T \mathbf{V}_H \mathbf{C}_\Delta$ , where:  $\hat{\phi}_H$  is the maximum likelihood estimate of  $\phi_H$ ; the matrix  $\mathbf{V}_H = \mathbf{J}_h(\hat{\phi}_S)^T \mathbf{V}_S \mathbf{J}_h(\hat{\phi}_S)$ ;  $\mathbf{J}_h(\hat{\phi}_S)$  is the Jacobian of the transformation  $\mathbf{h}(\phi_S)$ ; and  $\mathbf{V}_S$  is a blocked diagonal matrix consisting of the inverse observed information matrices for the separator nodes in each partition along the diagonal. The two posterior summaries  $\bar{\Delta}$  and  $\mathbf{S}_\Delta$ , namely the Bayes' estimator under a mean-squared error Bayes' risk function and corresponding variance-covariance matrix, may therefore be used under the general simultaneous inference framework of Hothorn et al. (2008); Bretz et al. (2011) to construct a multiplicity-adjusted test that  $\Delta = \mathbf{0}$ .

### 3.3 Simultaneous hypothesis testing

Given the estimator  $\bar{\Delta}$  and corresponding variance-covariance matrix  $\mathbf{S}_\Delta$ , define a vector of test statistics  $\mathbf{T}_n = \mathbf{D}_n^{-1/2}(\bar{\Delta} - \Delta)$ , where  $n$  is the dimension of the data  $\mathbf{y}$  and  $\mathbf{D}_n = \text{diag}(\mathbf{S}_\Delta)$ . Then it can be shown (Hothorn et al., 2008; Bretz et al., 2011) that  $\mathbf{T}_n$  tends in distribution to a multivariate normal distribution centred on  $\mathbf{0}$ :

$$\mathbf{T}_n \stackrel{a}{\sim} N_m(\mathbf{0}, \mathbf{R})$$

where  $\mathbf{R} := \mathbf{D}_n^{-1/2} \mathbf{S}_\Delta \mathbf{D}_n^{-1/2} \in \mathbb{R}^{m,m}$  is the posterior correlation matrix for the vector (length  $m$ ) of contrasts  $\Delta$ . Under the null hypothesis (3),  $\mathbf{T}_n = \mathbf{D}_n^{-1/2} \bar{\Delta} \stackrel{a}{\sim} N_m(\mathbf{0}, \mathbf{R})$ , and hence, assuming  $\mathbf{S}_\Delta$  is fixed and known, Hothorn et al. (2008); Bretz et al. (2011) show that a global  $\chi^2$ -test of conflict can be formulated:

$$X^2 = \mathbf{T}_n^T \mathbf{R}^+ \mathbf{T}_n \xrightarrow{d} \chi^2(\text{Rank}(\mathbf{R}))$$

where the authors use superscript  $+$  to denote the Moore-Penrose inverse of the corresponding matrix and  $\text{Rank}(\mathbf{R})$  is the degrees of freedom.

Importantly, it is also possible to construct multiply-adjusted local (individual) conflict tests, based on the  $m$   $z$ -scores corresponding to  $\mathbf{T}_n$  and the null distribution of the maximum



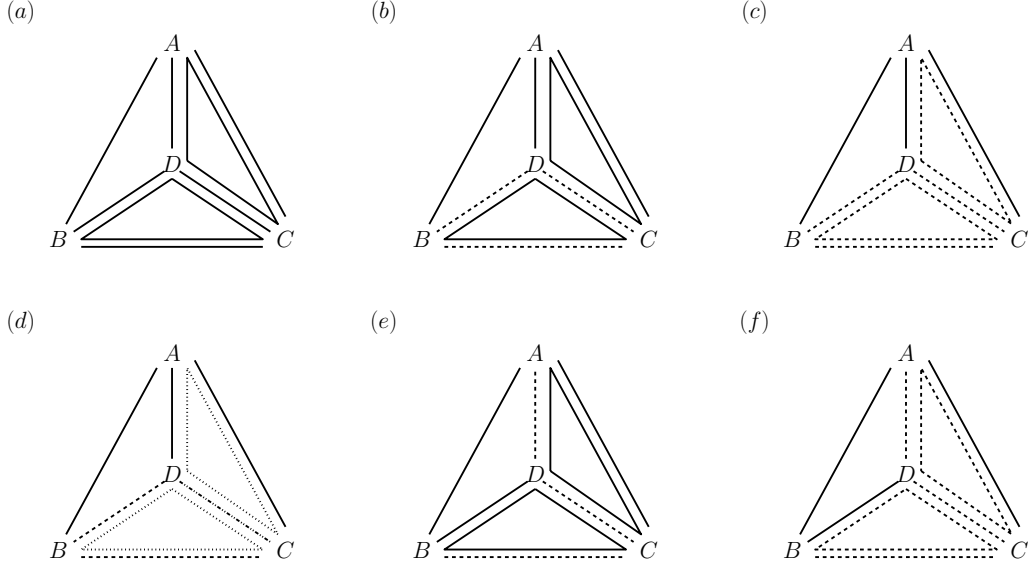


Figure 2: Smoking cessation evidence network, under (a) a consistency assumption; (b)-(f) inconsistency assumptions, where the evidence is partitioned in different ways. In (b), (c), (e) and (f), the direct evidence (dashed lines) is compared with the indirect evidence (solid lines) on each contrast where there is a dashed line. In (d), the evidence is separated into three spanning trees and a fourth partition for the multi-arm trial evidence.

of these,  $Z_{max}$ , (Hothorn et al., 2008; Bretz et al., 2011). This latter null distribution is obtained by integrating the limiting  $m$ -dimensional multivariate normal distribution over  $[-z, z]$  to obtain the cumulative distribution function  $\mathbb{P}(Z_{max} \leq z)$ . The individual conflict p-values are then calculated as  $\mathbb{P}(|z_k| < Z_{max}), k \in 1, \dots, m$ , with a corresponding global conflict p-value (an alternative to the  $\chi^2$ -test) given by  $\mathbb{P}(|z_{max}| < Z_{max})$ .

## 4 Examples

We now illustrate the idea of systematic multiple node-splitting to assess conflict in our two motivating examples. All analyses were carried out in **OpenBUGS** 3.2.2 (Lunn et al., 2009) and **R** 3.2.3 (R Core Team, 2015). We use the **R2OpenBUGS** package (Sturtz et al., 2005) to run **OpenBUGS** from within **R** and the **multcomp** package (Bretz et al., 2011) to carry out the simultaneous local and global max-T tests.

### 4.1 Network meta-analysis

Consider first a NMA in general, and for simplicity, assume initially there are no multi-arm trials and a common-effect model (equation (1)) for the data. The basic parameters  $\boldsymbol{\eta}_b$  form a spanning tree of the network of evidence (Figure 2), i.e. a graph with no cycles, such that each node in the network can be reached from every other node, either directly or indirectly through other nodes (van Valkenhoef et al., 2012). Multiple possible partitionings of the evidence network exist, so a choice has to be made (Figure 2). Suppose the spanning tree  $\boldsymbol{\eta}_b$

is identifiable by a set of evidence  $\mathbf{Y}_b$  containing outcomes from all trials designed to directly estimate the treatment effects in  $\boldsymbol{\eta}_b$ . Then every treatment effect is identifiable from  $\mathbf{Y}_b$ , by definition of a spanning tree and the fact that each treatment effect represented by edges outside the spanning tree is a functional parameter in the set  $\boldsymbol{\eta}_f$  that is a linear combination of the basic parameters. The data  $\mathbf{Y}_b$  therefore *indirectly* inform the functional parameters  $\boldsymbol{\eta}_f$ , whereas the remaining data,  $\mathbf{Y}_f = \mathbf{Y} \setminus \mathbf{Y}_b$  *directly* inform  $\boldsymbol{\eta}_f$ . A comparison between the direct and indirect evidence on  $\boldsymbol{\eta}_f$  can therefore be made, to assess conflict between the two types of evidence. The network is split into two partitions,  $\{\boldsymbol{\eta}_f^{Dir}, \mathbf{Y}_f\}$  (the “direct evidence partition”, DE) and  $\{\boldsymbol{\eta}_f^{Ind}, \mathbf{Y}_b\}$  (the “spanning tree partition”, ST) and the direct and indirect versions of the functional parameters compared:  $\boldsymbol{\Delta} = \boldsymbol{\eta}_f^{Dir} - \boldsymbol{\eta}_f^{Ind}$ . Note that depending on the studies that are in the DE partition, the basic parameters  $\boldsymbol{\eta}_b$  may also be (weakly) identifiable in the DE partition, due to prior information. Since a NMA model may be formulated as a DAG, this Direct/Indirect partitioning is equivalent to a multi-node split in the DAG at the functional parameters (Figure .2 of the Supplementary Material).

Generalising now to more complex situations, if the direct data  $\mathbf{Y}_f$  form a sub-network of evidence, the question arises of whether these data should be split into further partitions, by identifying a spanning tree for the sub-network. Then the vector  $\boldsymbol{\Delta}$  of contrasts to test would involve comparisons between more than two partitions, e.g. for three partitions:

$$\boldsymbol{\Delta} = (\boldsymbol{\eta}_f^1 - \boldsymbol{\eta}_f^2, \boldsymbol{\eta}_f^1 - \boldsymbol{\eta}_f^3, \boldsymbol{\eta}_f^2 - \boldsymbol{\eta}_f^3)^T$$

If we now consider a random rather than fixed heterogeneity effects model (equation (2)), a decision must be made on how to handle the variance components in  $\Sigma_\beta$ . One approach would be to split the variance components simultaneously with the means, so that  $\boldsymbol{\Delta}$  also includes contrasts for the variances. Alternatively, if the variance components are not well identified by the evidence in a partition, a common variance component could be assumed. Note that such commonality could potentially allow for feedback between partitions, since they would not be fully independent (Marshall and Spiegelhalter, 2007; Presanis et al., 2013).

Finally, for multi-arm trials, the key consideration is that multi-arm studies should have internal consistency, and hence their observations should not be split between partitions. A choice must therefore be made whether to initially include multi-arm data in the ST data  $\mathbf{Y}_b$ , in the DE data  $\mathbf{Y}_f$ , or in a third partition of their own. In the latter case, any study-specific treatment effect  $\mu_{di}^{JK}$ , where  $d$  is a multi-arm design, could be compared at least with the ST partition, where  $\eta^{JK}$  is definitely identified. Potentially, it could also be compared simultaneously with the DE partition, if the edge  $JK$  is identifiable in the DE partition. Note that the comparison can be made even if  $JK$  is not identifiable, or rather only weakly identifiable from the prior distribution, but if the prior is diffuse, then no conflict will be detected due to the uncertainty/lack of identifiability. Such a comparison is not therefore particularly meaningful, unless we are interested in prior-data conflict.

#### 4.1.1 Smoking cessation example

To illustrate concretely the above issues, we consider first the spanning tree  $(AB, AC, AD)$  corresponding to the parameters  $\boldsymbol{\eta}_b = \{\eta^{AB}, \eta^{AC}, \eta^{AD}\}$  for the smoking cessation example. Figures 2(b-d) demonstrate different ways of splitting the evidence based on this spanning

tree, depending on how we treat the evidence from multi-arm trials. In Figures 2(b,c), we consider just two partitions, with the multi-arm evidence either left in the ST partition  $\{\boldsymbol{\eta}_f^{Ind}, \mathbf{Y}_b\}$  or included in the DE partition  $\{\boldsymbol{\eta}_f^{Dir}, \mathbf{Y}_f\}$ , respectively. We compare the direct and indirect evidence on each of the edges or treatment comparisons  $(BC, BD, CD)$ . In Figure 2(d), we consider a series of spanning trees  $((AB, AC, AD), (BC, BD)$  and  $(CD))$ , together with a final partition consisting of evidence from multi-arm trials, resulting in four partitions.

We also consider an alternative choice of spanning tree,  $(AB, AC, BD)$ , as in Figures 2(e,f). In these two models, we again make a choice between including the multi-arm evidence in either the ST or DE partitions and compare the evidence in each partition on edges  $(AD, BC, CD)$ . In all cases, we assume random heterogeneity effects and make the choice to assume common variance components across the partitions, splitting only the means.

Table 1 gives posterior mean (sd) estimates of the treatment effects (log odds ratios) for edges outside the spanning tree, from each partition (1: spanning tree, 2: direct evidence for the two-partition models (b,c,e,f), and 1-3: sequential spanning trees, 4: multi-arm trial evidence, for the four-partition model (d)). Also given, for each edge outside the original spanning tree, are the posterior mean (sd) differences between partitions and both the local and global posterior probabilities of no difference, adjusted for the multiple tests and their correlation. First, note that the global test of no conflict varies by model, and hence by what partitions of evidence are compared with each other: the posterior probability of no conflict in model (b) is 94.7%, compared to only 23.4% and 27.4% for models (c) and (e). These latter two models appear to detect some evidence of conflict, despite the large uncertainty in many of the partition-specific treatment effect estimates, with several of the posterior standard deviations of the same order of magnitude as the corresponding posterior means, if not larger. Note also that the DIC is slightly smaller for the two models (c) and (e) which detect potential conflict, compared to those that don't.

A closer look at the local posterior probabilities of no conflict for each edge outside the initial spanning tree reveals that the potential conflict detected by models (c) and (e) involves edges including treatment  $D$  (posterior probabilities 17.8% and 18.6% for edges  $BD$  and  $CD$  respectively in model (c), 12.4% and 10.5% for edges  $AD$  and  $CD$  in model (e)). Each of these four local tests involves a partition where the estimated treatment effect for the relevant edge is implausibly large ( $> 6$  on the log odds ratio scale, i.e.  $> 400$  on the odds ratio scale). In model (e), the estimate in partition 2 for the  $AD$  edge of 12.7(6.2) is only directly informed by a single study, whose sample sizes are among the smallest in this network (study 7 in Table .1 of the Supplementary Material), with a 0 numerator for treatment  $A$ . In general, the studies included in partition 2 of model (e) (namely studies 7, 20, 23 and 24 in Table .1 of the Supplementary Material) are amongst the smallest sample sizes of the network. Note that by including the multi-arm studies in the spanning tree partition (Figure 2(e)), model (e) is equivalent to a model with a different spanning tree,  $(AD, CB, CD)$ , where the multi-arm trials are included in the partition outside of the spanning tree.

Unlike models (c), (e) and (f), where in both partitions, each sub-network spans all 4 treatments, in models (b) and (d), the spanning tree chosen,  $(AB, AC, AD)$ , is such that for each sub-network outside the spanning tree, not all the treatments are included (Figure 2). This results in a lack of identifiability for the basic parameters  $\boldsymbol{\eta}_b$  in partition 2 of model (b) and in partitions 2 and 3 of model (d) (Table 1), where their estimates are dominated by their

Table 1: Multiply adjusted posterior mean (sd) estimates of conflict between partitions, for each model (b)-(f) respectively. In the two-partition models (b,c,e,f), partition 1 is the spanning tree (indirect) evidence partition and partition 2 is the direct data partition. In model (d), partitions 1-3 are the sequential spanning trees and partition 4 is the multi-arm study partition.

ST: Model: Posterior:	(b)		AB,AC,AD (c)		(d)		AB,AC,BD (e)		(f)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
$AB_1$	0.472	(0.489)	0.338	(0.534)	0.329	( 0.568)	0.259	(0.429)	0.334	(0.566)
$AB_2$	-0.415	(5.276)	0.319	(0.983)	-0.230	( 5.849)	6.261	(3.251)	1.513	(1.041)
$AB_3$					-0.044	(10.009)				
$AB_4$					0.456	( 1.247)				
$\Delta_{AB_1-2}$										
$\Delta_{AB_1-3}$										
$\Delta_{AB_1-4}$										
$pAB_{1-2}$										
$pAB_{1-3}$										
$pAB_{1-4}$										
$AC_1$	0.877	(0.262)	0.814	(0.261)	0.828	( 0.280)	0.812	(0.238)	0.824	(0.272)
$AC_2$	-0.165	(5.262)	0.615	(0.866)	-0.379	( 5.848)	6.173	(3.140)	1.496	(0.835)
$AC_3$					0.114	( 7.051)				
$AC_4$					0.784	( 0.956)				
$\Delta_{AC_1-2}$										
$\Delta_{AC_1-3}$										
$\Delta_{AC_1-4}$										
$pAC_{1-2}$										
$pAC_{1-3}$										
$pAC_{1-4}$										
$AD_1$	1.010	(0.598)	9.337	(4.999)	9.508	( 5.330)	0.908	(0.794)	1.748	(1.526)
$AD_2$	0.262	(5.266)	0.679	(0.871)	0.871	( 5.859)	12.712	(6.225)	3.102	(1.690)
$AD_3$					0.319	( 7.044)				
$AD_4$					0.439	( 0.956)				
$\Delta_{AD_1-2}$							-11.804	(6.268)	-1.354	(2.279)
$\Delta_{AD_1-3}$										
$\Delta_{AD_1-4}$										
$pAD_{1-2}$							0.124		0.806	
$pAD_{1-3}$										
$pAD_{1-4}$										
$BC_1$	0.405	(0.527)	0.476	(0.590)	0.499	( 0.633)	0.553	(0.461)	0.490	(0.629)
$BC_2$	0.251	(0.808)	0.296	(0.577)	-0.149	( 1.015)	-0.087	(0.951)	-0.017	(0.693)
$BC_3$					0.158	(12.270)				
$BC_4$					0.329	( 0.957)				
$\Delta_{BC_1-2}$	0.155	(0.963)	0.180	(0.821)	0.649	( 1.198)	0.641	(1.059)	0.507	(0.937)
$\Delta_{BC_1-3}$					0.341	(12.290)				
$\Delta_{BC_1-4}$					0.171	( 1.161)				
$pBC_{1-2}$	0.986		0.971		0.979		0.807		0.839	
$pBC_{1-3}$					1.000					
$pBC_{1-4}$					1.000					
$BD_1$	0.538	(0.691)	8.999	(5.031)	9.180	( 5.357)	0.649	(0.530)	1.414	(1.173)
$BD_2$	0.678	(0.809)	0.360	(0.569)	1.101	( 1.026)	6.451	(3.083)	1.589	(0.802)
$BD_3$					0.363	(12.255)				
$BD_4$					-0.017	( 0.948)				
$\Delta_{BD_1-2}$	-0.140	(1.067)	8.639	(5.069)	8.079	( 5.444)				
$\Delta_{BD_1-3}$					8.817	(13.149)				
$\Delta_{BD_1-4}$					9.196	( 5.430)				
$pBD_{1-2}$	0.991		0.178		0.491					
$pBD_{1-3}$					0.952					
$pBD_{1-4}$					0.355					
$CD_1$	0.133	(0.594)	8.523	(5.011)	8.680	( 5.337)	0.095	(0.778)	0.924	(1.553)
$CD_2$	0.427	(0.680)	0.063	(0.460)	1.250	( 1.438)	6.539	(3.198)	1.606	(1.081)
$CD_3$					0.204	( 0.771)				
$CD_4$					-0.345	( 0.714)				
$\Delta_{CD_1-2}$	-0.294	(0.902)	8.459	(5.033)	7.430	( 5.526)	-6.443	(3.287)	-0.682	(1.893)
$\Delta_{CD_1-3}$					8.476	( 5.385)				
$\Delta_{CD_1-4}$					9.025	( 5.378)				
$pCD_{1-2}$	0.943		0.186		0.588		0.105		0.934	
$pCD_{1-3}$					0.430					
$pCD_{1-4}$					0.365					
Global p	0.947		0.234		0.200		0.274		0.733	
DIC	98.843		95.420		96.354		95.745		98.351	

diffuse prior distribution (Normal centred on 0 with standard deviation 10 on the log odds ratio scale). There is therefore no potential for detecting conflict about the basic parameters  $\eta_b$ , only about the functional parameters  $\eta_f$ .

The different results obtained from each of the five models are understandable, since each model partitions the evidence in a different way, and the detection of conflict relies on the conflicting evidence being in different rather than the same partitions. Note that where the same evidence is in the same partition for different models — for example, the evidence directly informing the  $AC$  edge in models (c) and (d) — approximately (within Monte Carlo error) the same estimate is reached in each model, as expected (0.81(0.26) in model (c), 0.82(0.28) in model (d), Table 1).

## 4.2 HIV prevalence evidence synthesis

Figure 3(b) demonstrates the multiple node-splits we make to systematically assess conflict in the original DAG of Figure 3(a), separating out the contributions of the prior model and each likelihood contribution. These node-splits result in 5 partitions, with 6 contrasts to test for equality to zero. Denoting the nodes in the “prior” partition (above the red arrows in Figure 3(b)) by the subscript  $p$  and the nodes in each “likelihood” partition (below the red arrows in Figure 3(b)) by  $d$ , the vector of contrasts to test is then

$$\Delta = (h(\rho_p) - h(\rho_d), h(\pi_p \kappa_p) - h([\pi \kappa]_d), h(\pi_p(1 - \kappa_p)) - h([\pi(1 - \kappa)]_d), \\ g(D_{L_p}) - g(D_{L_d}), g(D_{U_p}) - g(D_{U_d}), g(D_p) - g(D_d))^T$$

where  $h(\cdot)$  and  $g(\cdot)$  denote the logit and log functions respectively. These contrasts are represented by the red dot-dashed arrows in Figure 3(b).

In the “prior” partition, the priors given to the basic parameters are those of the original model (Section 2.2). In each “likelihood” partition, the basic parameters are given Jeffreys’ priors so that the posteriors in each of these partitions represent only the likelihood. These priors are Beta( $1/2, 1/2$ ) for the proportions and  $p(D_{B_d}) \propto 1/D_{B_d}^{1/2}$  for the lower and upper bounds ( $B = L, U$ ) for  $D$ .  $D_d$  is given a Uniform prior between  $D_{L_d}$  and  $D_{U_d}$ .

Figure 4 shows the posterior distributions of the contrasts  $\Delta$ , where 0 lies in these distributions and the corresponding unadjusted ( $p_U$ ) and multiply-adjusted ( $p_A$ ) individual conflict p-values testing for equality to 0. A global  $\chi^2$ -squared (Wald) test gives a conflict p-value of 0.001, suggesting conflict exists somewhere in the DAG. Examining the individual unadjusted (naive) conflict p-values would suggest prior-data conflict at the upper bound for the number diagnosed  $D_U$  (posterior probability the difference is zero is  $p_U = 0.008$ ) and hence the number diagnosed itself,  $D$  ( $p_U = 0.039$ ), as well as possibly at the proportion at risk,  $\rho$  ( $p_U = 0.078$ ). However, once the correlation between the individual tests has been taken into account, the posterior probabilities of the contrasts being zero (i.e. of no conflict) increase for all contrasts, albeit the probabilities are still low for  $D_U$  and  $D$ , at  $p_A = 0.175$  and  $p_A = 0.058$  respectively. Note that the posterior contrasts in Figure 4 are slightly non-normal, hence we interpret the adjusted posterior probabilities of no conflict as exploratory, rather than as absolute measures.

Examining closer the posterior distributions of the “prior” and “likelihood” versions of the node  $D$  (Figure .3, upper panel, in the Supplementary Material), we visualise better the




Figure 3: (a) DAG of initial model for synthesising Polish HIV prevalence data. (b) DAG of multiple node-split model comparing priors to each likelihood contribution. Note that the square brackets are used in denoting the nodes in the likelihood partition ( $[\pi\delta]_d, [\pi(1-\delta)]_d$ ) to emphasise the fact that these two nodes are independent parameters not functionally related to each other.




Figure 4: Posterior distributions of the contrasts  $\Delta$  for the HIV prevalence example. The red lines denote 0 difference,  $p_U$  is the unadjusted and  $p_A$  the multiply-adjusted individual conflict p-value respectively.

prior-data conflict: the “likelihood” version lies very much in the lower tail of the “prior” version. This is in spite of – or rather because of – the flat Uniform priors of the prior model, which translate into a non-Uniform implied prior for the function  $D_p = N\rho_p\pi_p\kappa_p$ .

The “saturated” model splitting apart each component of evidence in the DAG allows us to assess prior-data conflict in this model, but not conflict between different combinations of the likelihood evidence, due to lack of identifiability: in each likelihood partition in Figure 3(b), clearly only the parameter directly informed by the data, whether basic or functional, can be identified. To assess consistency of evidence between likelihood terms, we employ a cross-validators “leave-n-out” approach, for  $n = 1$  and  $n = 2$ , splitting in each case the relevant nodes *directly* informed by the left out data items. Note that other possibilities exist, to split for example at the basic parameters, depending on which data are left out. Table 2 gives unadjusted ( $p_U$ ) and various multiply-adjusted ( $p_{AW}, p_{AL}, p_{AA}$ ) individual posterior probabilities of no difference between nodes split between partitions 1 (the “left-out” evidence) and 2 (the remaining evidence). These posterior probabilities highlight inconsistency in the network of evidence  $\{y_1, y_2, y_4, y_5\}$ , i.e. informing the three nodes  $\rho, \pi\kappa$  and  $D = N\rho\pi\kappa$ . Note the low posterior probabilities of no difference concerning these three nodes in the “leave-1-out” models (A), (B) and (E), and in the “leave-2-out” models (B), (C), and (J) in particular. There is no potential for the evidence  $y_3$  on the prevalence of undiagnosed infection  $\pi(1 - \kappa)$  to conflict with any other evidence, since  $\pi$  and  $\kappa$  are not separately identifiable from the remaining evidence  $\{y_1, y_2, y_4, y_5\}$  alone. Hence all of the posterior probabilities of no difference concerning the node  $\pi(1 - \kappa)$  are high.

The conflict in the  $\{y_1, y_2, y_4, y_5\}$  network is well illustrated by the node-split model (J), where the count data on the lower and upper bounds for the  $D$  are “left out” in partition 1. Figure .3 (lower panel) of the Supplementary Material shows the posterior distributions for each of  $D_L, D_U$  and  $D$  in both partitions. Since in partition 2 the data on the limits for  $D$  have been excluded, the posterior distributions for the bounds (solid black and red lines) are flat and hugely variable. Despite this, the posterior distribution for  $D_2$  is relatively tightly peaked, due to the indirect evidence on  $D_2$  provided by the data informing  $\rho_2$  and  $\pi_2\kappa_2$ . It is this indirect evidence that conflicts with the direct evidence informing  $D_1$  via the data  $\{y_4, y_5\}$  on the bounds for  $D_1$ .

## 5 Discussion

We have proposed here the systematic assessment of conflict in an evidence synthesis, in particular accounting for the multiple tests for consistency entailed, through the simultaneous inference framework proposed by Hothorn et al. (2008); Bretz et al. (2011). We have chosen the max-T tests that allow both for multiply-adjusted local and global testing simultaneously.

Note that the use of this (typically classical) simultaneous inference framework relies on the asymptotic multivariate normality of the joint posterior distribution. In cases where the likelihood does not dominate the prior, resulting in a skewed or otherwise non-normal posterior, we treat the results of conflict analysis as exploratory, rather than absolute measures of conflict. If the posterior is skewed but still uni-modal, a global, implicitly multiply-adjusted, test for conflict can be formulated in terms of the Mahalanobis distance of each posterior sample from their mean, as we proposed in Presanis et al. (2013). This is a multivariate



Table 2: Results from “leave-n-out” node-split models for the Polish HIV data.  $p_U$  denotes the unadjusted conflict p-value;  $p_{AW}$  is the p-value adjusted for the multiple tests carried out *within* each model (A)-(J) for the leave-2-out approach;  $p_{AL}$  is the p-value adjusted for the 23 tests carried out in all models (A)-(J) for the leave-2-out approach; and  $p_{AA}$  is the p-value adjusted for 28 tests carried out in all leave-1-out models (A)-(E) and all leave-2-out models (A)-(J).

Model	Partition 1	Partition 2	Node split	$p_U$	$p_{AW}$	$p_{AL}$	$p_{AA}$
Leave-1-out							
(A)	$y_1$	$\{y_2, y_3, y_4, y_5\}$	$\rho$	$< 0.0001$		0.0060	0.0311
(B)	$y_2$	$\{y_1, y_3, y_4, y_5\}$	$\pi\kappa$	$< 0.0001$		0.0047	0.0246
(C)	$y_3$	$\{y_1, y_2, y_4, y_5\}$	$\pi(1 - \kappa)$	0.6201		0.9857	1.0000
(D)	$y_4$	$\{y_1, y_2, y_3, y_5\}$	$D_L$	0.1257		0.6242	0.9934
(E)	$y_5$	$\{y_1, y_2, y_3, y_4\}$	$D_U$	$< 0.0001$		0.5852	0.9890
Leave-2-out							
(A)	$\{y_1, y_2\}$	$\{y_3, y_4, y_5\}$	$\rho$	0.6972	0.7480	1.0000	1.0000
			$\pi\kappa$	0.2209	0.2230	0.9842	0.9937
(B)	$\{y_1, y_3\}$	$\{y_2, y_4, y_5\}$	$\rho$	$< 0.0001$	0.0023	0.0240	0.0294
			$\pi(1 - \kappa)$	0.4906	0.7717	1.0000	1.0000
(C)	$\{y_2, y_3\}$	$\{y_1, y_4, y_5\}$	$\pi\kappa$	$< 0.0001$	$< 0.0010$	$< 0.0010$	$< 0.0010$
			$\pi(1 - \kappa)$	0.8322	0.9000	1.0000	1.0000
			$\pi$	0.9921	0.9490	1.0000	1.0000
			$\kappa$	0.3329	0.6700	0.9998	1.0000
(D)	$\{y_1, y_4\}$	$\{y_2, y_3, y_5\}$	$\rho$	$< 0.0001$	0.0779	0.5754	0.6499
			$D_L$	0.0783	0.2851	0.9705	0.9866
(E)	$\{y_1, y_5\}$	$\{y_2, y_3, y_4\}$	$\rho$	0.0745	0.1260	0.7614	0.8271
			$D_U$	0.0026	0.0949	0.6543	0.7276
(F)	$\{y_2, y_4\}$	$\{y_1, y_3, y_5\}$	$\pi\kappa$	0.4682	0.9590	1.0000	1.0000
			$D_L$	0.0869	0.3000	0.9764	0.9898
(G)	$\{y_2, y_5\}$	$\{y_1, y_3, y_4\}$	$\pi\kappa$	0.4420	0.6690	1.0000	1.0000
			$D_U$	0.0137	0.1970	0.9044	0.9434
(H)	$\{y_3, y_4\}$	$\{y_1, y_2, y_5\}$	$\pi(1 - \kappa)$	0.1471	0.3330	0.9855	0.9944
			$D_L$	0.1328	0.3280	0.9844	0.9938
(I)	$\{y_3, y_5\}$	$\{y_1, y_2, y_4\}$	$\pi(1 - \kappa)$	0.5237	0.8100	1.0000	1.0000
			$D_U$	$< 0.0001$	0.2850	0.9702	0.9864
(J)	$\{y_4, y_5\}$	$\{y_1, y_2, y_3\}$	$D_L$	0.1958	0.5117	0.9933	0.9978
			$D_U$	$< 0.0001$	0.3963	0.9706	0.9866
			$D$	$< 0.0001$	0.0030	0.0213	0.0260

equivalent of calculating the tail area probability for regions further away from the posterior mean than the point  $\mathbf{0}$ . However, the Mahalanobis-based test does not allow us to obtain local tests for conflict, nor does it apply in the case of a multi-modal posterior. In the latter case, kernel density estimation could be used to obtain the multivariate tail area probability, although such estimation is computationally challenging for large posterior dimension.

Although generalised evidence syntheses have, for the most part, been carried out in a Bayesian framework, there are examples (e.g. Commenges and Hejblum, 2013) that are either frequentist or not fully Bayesian. And in the field of NMA, maximum likelihood and Bayesian methods are both employed (e.g. White et al., 2012; Jackson et al., 2014). An advantage of the simultaneous inference framework of Hothorn et al. (2008); Bretz et al. (2011) is that, given any estimator  $\bar{\Delta}$  of a vector of differences and its corresponding variance-covariance matrix  $\mathbf{S}_{\Delta}$ , regardless of the method used to obtain the estimator, the global and local max-T tests can be formulated.

Conflict p-values, both univariate and multivariate, can be seen as cross-validated posterior predictive checks (Presanis et al., 2013). There is a large literature on various types of Bayesian predictive diagnostics, including prior-, posterior- and mixed-predictive checks (e.g. Box, 1980; Gelman et al., 1996; Marshall and Spiegelhalter, 2007). A key issue much discussed in this literature is the lack of uniformity of posterior predictive p-values under the null hypothesis (Gelman, 2013), with such p-values conservative due to the double use of data. Much work has therefore been devoted to either alternative p-values (e.g. Bayarri and Berger, 2000) or post-processing of p-values to calibrate them (e.g. Steinbakk and Storvik, 2009). Gelman (2013) argues that the importance or otherwise of uniformity depends on the context in which the model checks are being made: in general non-uniformity is not an issue, but if the posterior predictive tests rely strongly on parameters or imputed latent data, then care should be taken. Since conflict p-values are cross-validated, the issue of conservatism and the double use of data does not apply, and in fact, for a wide class of standard hierarchical models, Gåsemyr (2016) has demonstrated the uniformity of the conflict p-value.

As illustrated by both applications, the choice of different ways of partitioning the evidence in a DAG can lead to different conclusions over the existence of conflict. This is to be expected when considering the local conflict p-values, since conflicting evidence may need to be in different partitions in order to be detectable. This is analogous to the idea of “masking” in cross-validated outlier detection, where outliers may not be detected if multiple outliers exist (Chaloner and Brant, 1988). In the case of the global tests for conflict, the NMA example showed that these are also not invariant to the choice of partition. In the NMA literature, alternative methods accounting for inconsistency include models that introduce “inconsistency parameters” that absorb any variability due to conflict beyond between-study heterogeneity (Lu and Ades, 2006; Higgins et al., 2012; Jackson et al., 2014). Higgins et al. (2012); Jackson et al. (2014) have pointed out that the apparent algorithm that Lu and Ades (2006) follow for identifying inconsistency parameters does not guarantee that all such parameters are identified, nor that the Lu-Ades model is invariant to the choice of baseline treatment. The authors further posit, and more recently have proved (Jackson et al., 2015), that their “design-by-treatment interaction model”, which introduces an inconsistency parameter systematically for each non-baseline treatment within each design, contains each possible Lu-Ades model as a sub-model. In related ongoing work, we note that each Lu-Ades model corresponds to a particular choice of node-splitting model, one being a reparameterisation of the other. The

lack of invariance of results of testing for inconsistency from one Lu-Ades model to another is therefore not surprising, since, as we illustrated here, different choices of node-splitting model correspond to different partitions of evidence being compared. The lack of invariance of a global test for conflict to the choice of node-splitting model, although unsurprising, is perhaps unsatisfactory: however, as we illustrated in this paper, this lack clearly emphasises the need for a more comprehensive and systematic assessment of conflict throughout a DAG, both at a local level and across different types of node-split model, than just a single global test can provide. The design-by-treatment interaction model provides a way of doing so and we are further investigating the relationship of the (fixed inconsistency effects) design-by-treatment interaction model to a “saturated” node-splitting model where each edge in the network is split.

Note that in the NMA example considered here, we have concentrated on a “contrast-based” as opposed to “arm-based” parameterisation (Hong et al., 2016; Dias and Ades, 2016). Note also that we have considered the case where each study has a study-specific baseline treatment  $B_d$  and the network as a whole has a baseline treatment  $A$ . However, alternative parameterisations could be considered, such as using a two-way linear predictor with main effects for both treatment and study, treating the counter-factual or missing treatment designs as missing data (Jones et al., 2011; Piepho et al., 2012). Although we have not yet explored alternative parameterisations, we posit that systematic node-splitting could be applied equally in such alternatives.

As with any cross-validators work, the systematic assessment of conflict at every node in a DAG can quickly become computationally burdensome as a model grows in dimension. An area for future research is the systematic analysis of conflict using efficient algorithms (Lunn et al., 2013; Goudie et al., 2015) in a Markov melding framework (Goudie et al., 2016) which allows for an efficient modular approach to model building.

## Acknowledgements

This work was supported by the Medical Research Council [Unit Programme number U105260566]; and the Polish National Science Centre [grant no. DEC-2012/05/E/ST1/02218]. The authors also thank Ian White and Dan Jackson for their very helpful comments.

## References

- Ades, A. E. and Sutton, A. J. (2006). Multiparameter evidence synthesis in epidemiology and medical decision-making: current approaches. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **169**, 5–35.
- Bayarri, M. J. and Berger, J. O. (2000). P Values for Composite Null Models. *Journal of the American Statistical Association* **95**, 1127–1142.
- Bernardo, J. M. and Smith, A. F. M. (1994). *Bayesian Theory*. John Wiley & Sons, Inc., Hoboken, NJ, USA.

- Box, G. E. P. (1980). Sampling and Bayes' Inference in Scientific Modelling and Robustness. *Journal of the Royal Statistical Society. Series A (General)* **143**, 383–430.
- Bretz, F., Hothorn, T., and Westfall, P. (2011). *Multiple Comparisons Using R*. Chapman and Hall/CRC, first edition.
- Chaloner, K. and Brant, R. (1988). A bayesian approach to outlier detection and residual analysis. *Biometrika* **75**, 651–659.
- Commenges, D. and Hejblum, B. (2013). Evidence synthesis through a degradation model applied to myocardial infarction. *Lifetime Data Analysis* **19**, 1–18.
- De Angelis, D., Presanis, A. M., Birrell, P. J., Tomba, G. S., and House, T. (2014). Four key challenges in infectious disease modelling using data from multiple sources. *Epidemics*.
- Dias, S. and Ades, A. E. (2016). Absolute or relative effects? arm-based synthesis of trial data. *Research Synthesis Methods* **7**, 23–28. RSM-07-2015-0036.R1.
- Dias, S., Welton, N. J., Caldwell, D. M., and Ades, A. E. (2010). Checking consistency in mixed treatment comparison meta-analysis. *Statistics in Medicine* **29**, 932–944.
- Gelman, A. (2013). Two simple examples for understanding posterior p-values whose distributions are far from uniform. *Electronic Journal of Statistics* **7**, 2595–2602.
- Gelman, A., Meng, X.-L., and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* **6**, 733–807.
- Goudie, R. J. B., Hovorka, R., Murphy, H. R., and Lunn, D. (2015). Rapid model exploration for complex hierarchical data: application to pharmacokinetics of insulin aspart. *Statistics in Medicine* **34**, 3144–3158.
- Goudie, R. J. B., Presanis, A. M., Lunn, D. J., De Angelis, D., and Wernisch, L. (2016). Model surgery: joining and splitting models with Markov melding. <https://arxiv.org/abs/1607.06779>.
- Gåsemýr, J. (2016). Uniformity of Node Level Conflict Measures in Bayesian Hierarchical Models Based on Directed Acyclic Graphs. *Scandinavian Journal of Statistics* **43**, 20–34.
- Gåsemýr, J. and Natvig, B. (2009). Extensions of a conflict measure of inconsistencies in Bayesian hierarchical models. *Scandinavian Journal of Statistics* **36**, 822–838.
- Higgins, J. P. T., Jackson, D., Barrett, J. K., Lu, G., Ades, A. E., and White, I. R. (2012). Consistency and inconsistency in network meta-analysis: concepts and models for multi-arm studies. *Research Synthesis Methods* **3**, 98–110.
- Hong, H., Chu, H., Zhang, J., and Carlin, B. P. (2016). A bayesian missing data framework for generalized multiple outcome mixed treatment comparisons. *Research Synthesis Methods* **7**, 6–22. RSM-01-2014-0002.R2.
- Hothorn, T., Bretz, F., and Westfall, P. (2008). Simultaneous inference in general parametric models. *Biometrical Journal* **50**, 346–363.
- Jackson, D., Barrett, J. K., Rice, S., White, I. R., and Higgins, J. P. T. (2014). A design-by-treatment interaction model for network meta-analysis with random inconsistency effects. *Statistics in Medicine* **33**, 3639–3654.


- Jackson, D., Boddington, P., and White, I. R. (2015). The design-by-treatment interaction model: a unifying framework for modelling loop inconsistency in network meta-analysis. *Research Synthesis Methods* **7**, 329–32.
- Jones, B., Roger, J., Lane, P. W., Lawton, A., Fletcher, C., Cappelleri, J. C., Tate, H., and Moneuse, P. (2011). Statistical approaches for conducting network meta-analysis in drug development. *Pharmaceutical Statistics* **10**, 523–531.
- Krahn, U., Binder, H., and König, J. (2013). A graphical tool for locating inconsistency in network meta-analyses. *BMC Medical Research Methodology* **13**, 35+.
- Lauritzen, S. L. (1996). *Graphical Models*. Oxford Statistical Science Series. Oxford University Press, USA.
- Lu, G. and Ades, A. E. (2006). Assessing Evidence Inconsistency in Mixed Treatment Comparisons. *Journal of the American Statistical Association* **101**, 447–459.
- Lumley, T. (2002). Network meta-analysis for indirect treatment comparisons. *Statistics in Medicine* **21**, 2313–2324.
- Lunn, D., Barrett, J., Sweeting, M., and Thompson, S. (2013). Fully Bayesian hierarchical modelling in two stages, with application to meta-analysis. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **62**, 551–572.
- Lunn, D., Spiegelhalter, D., Thomas, A., and Best, N. (2009). The BUGS project: Evolution, critique and future directions. *Statistics in Medicine* **28**, 3049–3067.
- Marshall, E. C. and Spiegelhalter, D. J. (2007). Identifying outliers in Bayesian hierarchical models: a simulation-based approach. *Bayesian Analysis* **2**, 409–444.
- Piepho, H. P., Williams, E. R., and Madden, L. V. (2012). The use of Two-Way linear mixed models in multitreatment Meta-Analysis. *Biometrics* **68**, 1269–1277.
- Presanis, A. M., Ohlssen, D., Spiegelhalter, D. J., and De Angelis, D. (2013). Conflict diagnostics in directed acyclic graphs, with applications in Bayesian evidence synthesis. *Statistical Science* **28**, 376–397.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Rosinska, M., Gwiazda, P., Angelis, D. D., and Presanis, A. M. (2016). Bayesian evidence synthesis to estimate HIV prevalence in men who have sex with men in Poland at the end of 2009. *Epidemiology & Infection* **144**, 1175–1191.
- Salanti, G. (2012). Indirect and mixed-treatment comparison, network, or multiple-treatments meta-analysis: many names, many benefits, many concerns for the next generation evidence synthesis tool. *Research Synthesis Methods* **3**, 80–97.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and van der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64**, 583–639.
- Steinbakk, G. H. and Storvik, G. O. (2009). Posterior Predictive p-values in Bayesian Hierarchical Models. *Scandinavian Journal of Statistics* **36**, 320–336.

- Sturtz, S., Ligges, U., and Gelman, A. (2005). R2WinBUGS: A Package for Running WinBUGS from R. *Journal of Statistical Software* **12**, 1–16.
- van Valkenhoef, G., Dias, S., Ades, A. E., and Welton, N. J. (2016). Automated generation of node-splitting models for assessment of inconsistency in network meta-analysis. *Res. Syn. Meth.* **7**, 80–93.
- van Valkenhoef, G., Tervonen, T., de Brock, B., and Hillege, H. (2012). Algorithmic parameterization of mixed treatment comparisons. *Statistics and Computing* **22**, 1099–1111.
- Welton, N. J., Sutton, A. J., Cooper, N. J., Abrams, K. R., and Ades, A. E. (2012). *Evidence Synthesis in a Decision Modelling Framework*. John Wiley & Sons, Ltd.
- White, I. R., Barrett, J. K., Jackson, D., and Higgins, J. P. T. (2012). Consistency and inconsistency in network meta-analysis: model estimation using multivariate meta-regression. *Research Synthesis Methods* **3**, 111–125.

## Supplementary Material

### Appendix A: Figures and Tables

(a)



(b)

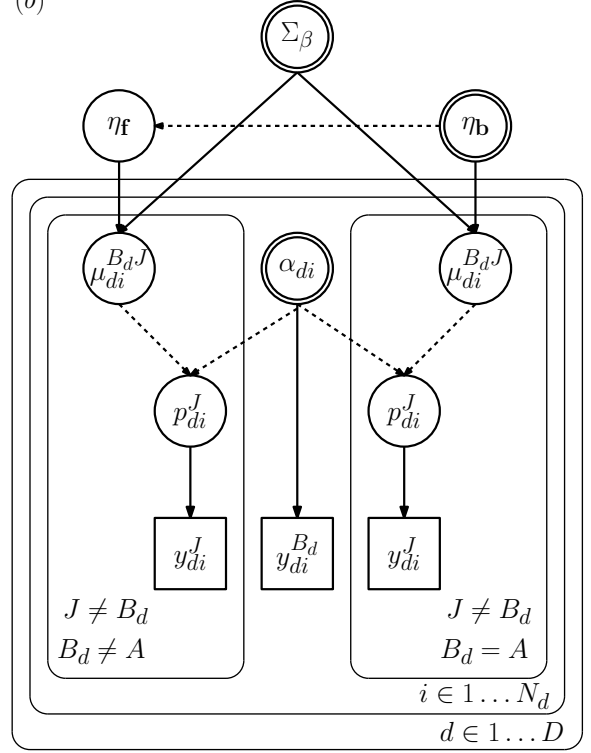


Figure .1: (a) DAG of NMA under assumptions of a common treatment effect  $\eta^{JK}$  (no heterogeneity) and consistency  $\eta^{JK} = \eta^{AK} - \eta^{AJ}$ . (b) DAG of NMA under assumptions of random treatment effects, to account for heterogeneity, and consistency.

Table .1: Smoking cessation data set

Study	Design	$y_A$	$n_A$	$y_A/n_A$	$y_B$	$n_B$	$y_B/n_B$	$y_C$	$n_C$	$y_C/n_C$	$y_D$	$n_D$	$y_D/n_D$
1	AB	79	702	0.113	77	694	0.111	.	.	.	.	.	.
2	AB	18	671	0.027	21	535	0.039	.	.	.	.	.	.
3	AB	8	116	0.069	19	149	0.128	.	.	.	.	.	.
4	AC	75	731	0.103	.	.	.	363	714	0.508	.	.	.
5	AC	2	106	0.019	.	.	.	9	205	0.044	.	.	.
6	AC	58	549	0.106	.	.	.	237	1561	0.152	.	.	.
7	AC	0	33	0.000	.	.	.	9	48	0.188	.	.	.
8	AC	3	100	0.030	.	.	.	31	98	0.316	.	.	.
9	AC	1	31	0.032	.	.	.	26	95	0.274	.	.	.
10	AC	6	39	0.154	.	.	.	17	77	0.221	.	.	.
11	AC	64	642	0.100	.	.	.	107	761	0.141	.	.	.
12	AC	5	62	0.081	.	.	.	8	90	0.089	.	.	.
13	AC	20	234	0.085	.	.	.	34	237	0.143	.	.	.
14	AC	95	1107	0.086	.	.	.	143	1031	0.139	.	.	.
15	AC	15	187	0.080	.	.	.	36	504	0.071	.	.	.
16	AC	78	584	0.134	.	.	.	73	675	0.108	.	.	.
17	AC	69	1177	0.059	.	.	.	54	888	0.061	.	.	.
18	ACD	9	140	0.064	.	.	.	23	140	0.164	10	138	0.072
19	AD	0	20	0.000	.	.	.	.	.	.	9	20	0.450
20	BC	.	.	.	20	49	0.408	16	43	0.372	.	.	.
21	BCD	.	.	.	11	78	0.141	12	85	0.141	29	170	0.171
22	BD	.	.	.	7	66	0.106	.	.	.	32	127	0.252
23	CD	.	.	.	.	.	.	12	76	0.158	20	74	0.270
24	CD	.	.	.	.	.	.	9	55	0.164	3	26	0.115



Table .2: Treatment effect posterior estimates (mean (sd)) for the common- and random-effect models respectively, with deviance summaries: posterior mean deviance  $\mathbb{E}_{\theta|y}(D)$ ; deviance evaluated at posterior means  $D(\mathbb{E}_{\theta|y}\theta)$ ; effective number of parameters  $p_D$ ; and deviance information criterion  $DIC$ .

Model: $\mu^{JK}$ :	Common-effect		Random-effect	
	Posterior mean	Posterior sd	Posterior mean	Posterior sd
$AB$	0.224	(0.124)	0.496	(0.405)
$AC$	0.765	(0.059)	0.843	(0.236)
$AD$	0.840	(0.174)	1.103	(0.439)
$BC$	0.541	(0.132)	0.347	(0.419)
$BD$	0.616	(0.192)	0.607	(0.492)
$CD$	0.075	(0.171)	0.260	(0.418)
$\mathbb{E}_{\theta y}(D)$	267		54	
$D(\mathbb{E}_{\theta y}\theta)$	240		10	
$p_D$	27		44	
$DIC$	294		98	

Table .3: Results from initial HIV model: observations; posterior mean (sd) estimates; posterior mean deviance  $\mathbb{E}_{\theta|y}(D)$ ; deviance evaluated at posterior means  $D(\mathbb{E}_{\theta|y}\theta)$ ; effective number of parameters  $p_D$ ; and deviance information criterion  $DIC$ .

Parameter	Data			Estimates		Deviance summaries			
$\theta$	$y$	$n$	$y/n$	$\hat{y}$	$\hat{\theta}$	$\mathbb{E}_{\theta y}(D)$	$D(\mathbb{E}_{\theta y}\theta)$	$p_D$	$DIC$
$\rho$	35	1536	0.023	14.6 ( 1.5)	0.010 (0.001)	21.0	20.7	0.4	21.4
$\pi\kappa$	113	2840	0.040	92.5 ( 8.9)	0.033 (0.003)	5.5	4.4	1.1	6.5
$\pi(1-\kappa)$	136	2725	0.050	136.7 (11.3)	0.050 (0.004)	1.0	0.0	1.0	2.0
$D_L$	836			836.2 (28.9)	836.2 (28.9)	1.0	0.0	1.0	2.0
$D_U$	5034			5054.3 (70.8)	5054.4 (70.8)	1.1	0.1	1.0	2.1
Total						29.5	25.1	4.4	33.9




Figure .2: DAG of common-effect network meta-analysis model, split into direct (DE) and indirect (ST) evidence informing the functional parameters  $\eta_f$ , i.e. those edges outside of the spanning tree formed by the basic parameters  $\eta_b$ .

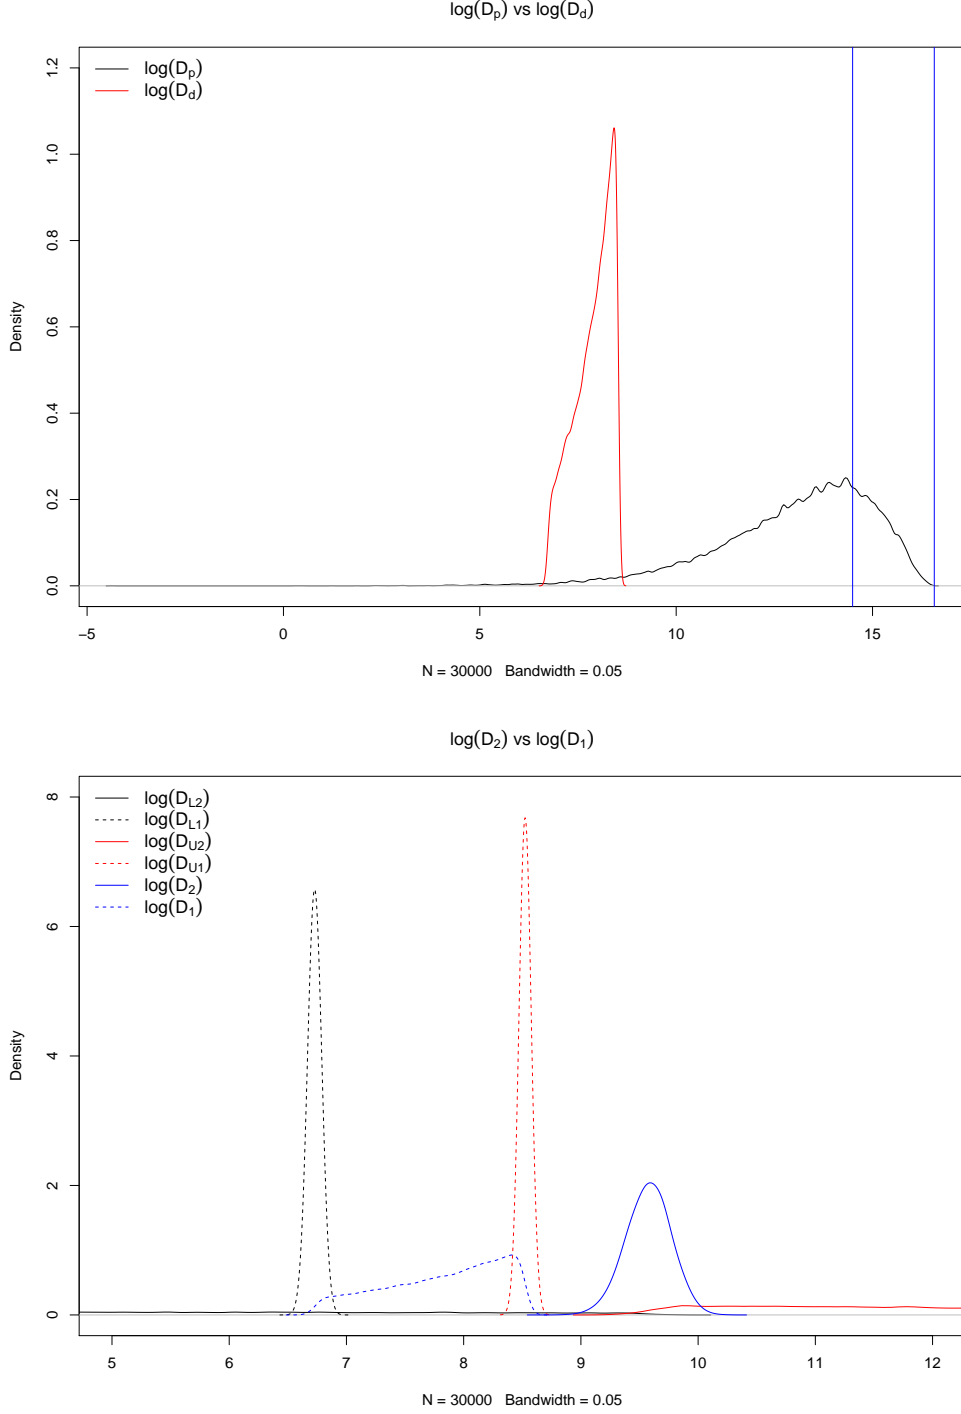


Figure 3: Upper panel: Posterior distributions of the nodes  $D_p$  and  $D_d$  for the HIV prevalence example, on the log scale. The right-hand blue line denotes where the total population of Poland ( $N = 15,749,944$ ) lies, i.e. the maximum possible value *a priori* for the number diagnosed. The left-hand blue line denotes the value  $\log(N \times 0.5^3)$ , i.e. the prior mean of  $\log(D_p) = \log(N\rho_p\pi_p\kappa_p)$ . Lower panel: Posterior distributions of the nodes  $D_{L1}, D_{L2}, D_{U1}, D_{U2}, D_1$  and  $D_2$  for the HIV prevalence “leave-2-out” node-split model (J), on the log scale. The dashed lines represent the nodes in partition 1, i.e. the “left-out” partition, where the posteriors are based only on the likelihood given by  $\{y_4, y_5\}$  and Jeffreys’ priors for  $D_{L1}, D_{U1}$ . The solid lines give the corresponding posteriors in partition 2, i.e. based on all the original model priors and on the dataset  $\{y_1, y_2, y_3\}$ .

## Appendix B: Asymptotics

Let  $p(\boldsymbol{\theta}_1), \dots, p(\boldsymbol{\theta}_Q)$  denote the set of prior distributions for the basic parameters  $\boldsymbol{\theta}_q$  in each partition  $q$ . Then by the independence of each partition, the joint posterior distribution of all parameters  $\boldsymbol{\phi}$  in all partitions is

$$p(\boldsymbol{\phi} \mid \mathbf{y}) = \prod_{q=1}^Q p(\boldsymbol{\theta}_q) p(\mathbf{y}_q \mid \boldsymbol{\theta}_q).$$

If the joint prior distribution is dominated by the likelihood, then asymptotically (Bernardo and Smith, 1994), the joint posterior distribution of all nodes is multi-variate normal:

$$\boldsymbol{\phi} \mid \mathbf{y} \stackrel{a}{\sim} N_{\sum_q n_q} \left( (\hat{\boldsymbol{\phi}}_1, \dots, \hat{\boldsymbol{\phi}}_Q), \mathbf{V} \right)$$

where  $n_q$  is the total number of parameters in partition  $q$ , whether basic or not, and  $\mathbf{V}$  is the inverse observed information matrix for the parameters  $\boldsymbol{\phi}$ . Since the vector of separator nodes,  $\boldsymbol{\phi}_S = (\boldsymbol{\phi}_1^{(s)}, \dots, \boldsymbol{\phi}_Q^{(s)})$ , is a subset of  $\boldsymbol{\phi}$ , their joint posterior is also multivariate normal:

$$\boldsymbol{\phi}_S \mid \mathbf{y} \stackrel{a}{\sim} N_m \left( (\hat{\boldsymbol{\phi}}_1^{(s)}, \dots, \hat{\boldsymbol{\phi}}_Q^{(s)}), \mathbf{V}_S \right) \quad (4)$$

where  $m = \sum_q m_q$  is the total number of separator nodes, including node-split copies, and  $\mathbf{V}_S$  is the appropriate sub-matrix of  $\mathbf{V}$ . Since the partitions are independent,  $\mathbf{V}_S$  is a blocked diagonal matrix consisting of the inverse observed information matrices for separator nodes in each partition along the diagonal.

By theorem 5.17 of Bernardo and Smith (1994), since (4) holds and if  $\mathbf{J}_h(\boldsymbol{\phi}_S) = \frac{\partial \mathbf{h}(\boldsymbol{\phi}_S)}{\partial \boldsymbol{\phi}_S}$  is non-singular with continuous entries, then the posterior distribution of the transformed separator nodes,  $\boldsymbol{\phi}_H = \mathbf{h}(\boldsymbol{\phi}_S)$ , is also asymptotically normal:

$$\boldsymbol{\phi}_H \mid \mathbf{y} \stackrel{a}{\sim} N_m \left( \mathbf{h}(\hat{\boldsymbol{\phi}}_1^{(s)}, \dots, \hat{\boldsymbol{\phi}}_Q^{(s)}), \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S)^T \mathbf{V}_S \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S) \right)$$

The Jacobian  $\mathbf{J}_h(\boldsymbol{\phi}_S)$  exists and is non-singular for the sorts of transformations we use in practice, for example log and logit transformations.

A further application of theorem 5.17 of Bernardo and Smith (1994) results in a posterior distribution of the contrasts  $\boldsymbol{\Delta}$  that is also asymptotically multivariate normal, if  $\frac{\partial \boldsymbol{\Delta}(\boldsymbol{\phi})}{\partial \boldsymbol{\phi}} = \mathbf{C}_\Delta^T$  is non-singular with continuous entries, which as a contrast matrix it is:

$$\begin{aligned} \boldsymbol{\Delta} \mid \mathbf{y} &\stackrel{a}{\sim} N_p \left( \mathbf{C}_\Delta^T \mathbf{h}(\hat{\boldsymbol{\phi}}_1^{(s)}, \dots, \hat{\boldsymbol{\phi}}_Q^{(s)}), \mathbf{C}_\Delta^T \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S)^T \mathbf{V}_S \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S) \mathbf{C}_\Delta \right) \\ &= N_p \left( \mathbf{C}_\Delta^T \hat{\boldsymbol{\phi}}_H, \mathbf{C}_\Delta^T \mathbf{V}_H \mathbf{C}_\Delta \right) \end{aligned} \quad (5)$$

for  $\mathbf{V}_H = \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S)^T \mathbf{V}_S \mathbf{J}_h(\hat{\boldsymbol{\phi}}_S)$ . Asymptotically, therefore, the posterior mean  $\overline{\boldsymbol{\Delta}} = \mathbf{C}_\Delta^T \overline{\boldsymbol{\phi}}_H \stackrel{a}{\approx} \mathbf{C}_\Delta^T \hat{\boldsymbol{\phi}}_H$  and the posterior variance-covariance matrix of  $\boldsymbol{\Delta}$  is  $\mathbf{S}_\Delta \stackrel{a}{\approx} \mathbf{C}_\Delta^T \mathbf{V}_H \mathbf{C}_\Delta$ .